



Fog Computing

Tobias Pfandzelter | Part 6: Low-Earth Orbit Edge Computing

Agenda

Lectures

From Cloud to
Fog Computing

Data
Distribution

Data
Management

Platforms &
Applications

Testing &
Benchmarking

LEO Edge
Computing

Assignments

Prototyping
Assignment

Reading
Assignment

Wrap-up

Q&A


Final Test

Low-Earth Orbit Edge Computing


All the stuff you have already learned about. But in space.

Context

Types of Edge Computing in Orbit




- In-network computing
- Piggy-backing LEO satellite broadband networks
- Same applications as terrestrial fog computing
- Large scale, highly dynamic




- Edge, on-device computing
- Processing to lower downlink bandwidth and energy consumption
- On-satellite sensor data (e.g., imaging)
- (Mostly) isolated

Types of Edge Computing in Orbit




- In-network computing
- Piggy-backing LEO satellite broadband networks
- Same applications as terrestrial fog computing
- Large scale, highly dynamic



- Edge, on-device computing
- Processing to lower downlink bandwidth and energy consumption
- On-satellite sensor data (e.g., imaging)
- (Mostly) isolated

In-Network LEO Edge Computing

Internet from Space



[14] P. Kirstein, "University College London ARPANET Project, Annual Report 1977," Apr. 1978.

LEO Satellites

NASA awards millions to SpaceX and Amazon's Project Kuiper for satellite communications

BY ALAN BOYLE on April 20, 2022 at 3:24 pm

<https://www.geekwire.com/2022/nasa-awards-millions-to-spacex-and-amazons-project-kuiper-for-satellite-communications/>

Arianespace Soyuz rocket launches 34 OneWeb internet satellites into space

By Mike Wall published August 22, 2021

Arianespace has now lofted a total of 288 OneWeb satellites.



<https://www.space.com/arianespace-soyuz-launches-oneweb-9-internet-satellite-mission>

INVESTING IN SPACE

About 150,000 people in Ukraine are using SpaceX's Starlink internet service daily, government official says


PUBLISHED MON, MAY 2 2022-12:03 PM EDT | UPDATED MON, MAY 2 2022-1:25 PM EDT

<https://www.cnn.com/2022/05/02/ukraine-official-150000-using-spacexs-starlink-daily.html>

SpaceX deploys more Starlink satellites, aims for higher launch cadence

© April 21, 2022 Stephen Clark


<https://spaceflightnow.com/2022/04/21/spacex-deploys-more-starlink-satellites-aims-for-higher-launch-cadence/>




[1]




GEO Satellites



LEO Satellites



LEO Satellite Mobility




LEO Satellite Mobility






550km Orbit → 95 minute orbital period


LEO Satellite Mobility



LEO Satellite Mobility



Topologies





Starlink Phase I Constellation

- 4,409 satellites in five “shells”
- Walker Delta constellation
- Each shell is independent



Altitude in km	Planes	Nodes per Plane	Orbital Inclination	Color
550	72	22	53°	blue
1,110	32	50	53.8°	orange
1,130	8	50	74°	purple
1,275	5	75	81°	red
1,325	6	75	70°	green

[18] G. Karabulut Kurt, “An Overview of Communications and Networking in LEO Mega-Constellations,” Working in Space Together: IEEE Low-Earth Orbit Satellite Systems Workshop, Nov. 2022.





Networking




Networking: Delay Advantage

Impact of physical attributes that constitute the first order effects:


- Altitude
- Speed of light: **c** vs. **0.68c**
- Distance
- Constellation Density

Note: Despite their importance, routing, forwarding, queueing, etc. are considered 2nd or 3rd order effects




Stolen from: [20] A. Akhavain, “Non Terrestrial Networks: Laying out the foundation for ubiquitous network coverage” in LEOCONN Webinar Series, Apr. 2023.

LEO Satellite Edge Computing




LEO Satellite Edge Computing




[2] D. Bhattacharjee, S. Kassing, M. Licciardello, and A. Singla, "In-orbit computing: An outlandish thought experiment?" in Proceedings of the 19th ACM Workshop on Hot Topics in Networks (HotNets 2019), Nov. 2020, pp. 197–204.

LEO Satellite Edge Computing




In-orbit compute is accessible from everywhere at low latency. The dotted lines are maximum latency i.e., to the farthest reachable satellite-server, while the solid lines are minimum latency i.e., to the nearest satellite-server



A substantial number of satellite-servers would be reachable from everywhere at all times. The solid lines show the average number across time, with the range over time shown by dotted and dashed lines. The peak at certain latitudes (less pronounced for Kuiper) arises from orbital geometry of the planned constellations.

[2] D. Bhattacharjee, S. Kassing, M. Licciardello, and A. Singla, "In-orbit computing: An outlandish thought experiment?" in Proceedings of the 19th ACM Workshop on Hot Topics in Networks (HotNets 2019), Nov. 2020, pp. 197–204.


LEO Edge: CDNs



[3] T. Pfandzelter et al., “Edge (of the Earth) Replication: Optimizing Content Delivery in Large LEO Satellite Communication Networks” in Proceedings of the 21st IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGrid 2021), May. 2021, pp. 565–575.


[4] U. Speidel, “LEO mega-constellations - where the challenges lie” in LEOCONN 2022, Oct. 2022.

LEO Edge: CDNs




[19] Akamai Technologies, "Object delivery,"
https://developer.akamai.com/legacy/learn/Content_Delivery/Object_Delivery.html 2019.

LEO Edge: Remote Internet of Things




LEO Edge: Remote Internet of Things




[12] NOAA Center for Tsunami Research, "DART 4G: Deep-ocean Assessment and Reporting of Tsunami - 4th Generation Tsunami Measurement System," May. 2016.

LEO Edge: Multi-User Interaction



[2] D. Bhattacharjee, S. Kassing, M. Licciardello, and A. Singla, “In-orbit computing: An outlandish thought experiment?” in Proceedings of the 19th ACM Workshop on Hot Topics in Networks (HotNets 2019), Nov. 2020, pp. 197–204.

LEO Edge: Multi-User Interaction



[2] D. Bhattacharjee, S. Kassing, M. Licciardello, and A. Singla, “In-orbit computing: An outlandish thought experiment?” in Proceedings of the 19th ACM Workshop on Hot Topics in Networks (HotNets 2019), Nov. 2020, pp. 197–204.

Isn't this way too expensive???

- **Weight and volume:** HPE ProLiant DL325 Gen10 = 15.6kg, ca. 6% weight of a Starlink satellite
- **Radiation hardening:** not difficult (we'll get to that)
- **Power:** Starlink satellite solar output is estimated at 1.5kW, server needs 225W-350W (15-23%)
- **Heat:** Probably radiators or thermoelectric harvesters required
- **Life-cycle:** ~5 years for a Starlink satellite, 3 years for a data center server
- **Cost:**
 - ~42,000 USD server launch cost
 - Data center server yearly TCO 5,000USD
 - At 3-year lifespan, satellite-server only roughly 3x as expensive as data center server

[2] D. Bhattacharjee, S. Kassing, M. Licciardello, and A. Singla, "In-orbit computing: An outlandish thought experiment?" in Proceedings of the 19th ACM Workshop on Hot Topics in Networks (HotNets 2019), Nov. 2020, pp. 197–204.

Failure Scenarios

- Up-/Downlink Degradation: ca. 50% bandwidth reduction during heavy rain (at *any* ground station)
- On-board Compute Failure: with 1mm aluminum shielding TID is below TID limit of Snapdragon SA8155P (automotive processor) for a five-year lifetime at 550km altitude → SEU-induced soft error rate on the order of 10^{-3} to 10^{-4} per day
- Orbital Maneuvers: little to no impact on overall network characteristics
- ISL degradation: not an issue with lasers in space
- Adversarial attacks: always possible

[17] T. Pfandzelter and D. Bermbach, “Failure is not an Option: Considerations for Software Fault-Tolerance in LEO Satellite Edge Computing”, Feb. 2023, arXiv:2302.08952.

Compute Platforms and Abstractions


Easy Enough?

Consider...


- **Mobility:** 90-minute orbital period
- **Coverage:** satellite handover every 1-2 minutes
- **Resource Constraints:** a server in space?
- **Fault Tolerance:** *a server in space???*
- **Multi-Tenancy:** different applications for different geographies
- **Scalability:** thousands of satellites

[5] T. Pfandzelter et al., “Towards a Computing Platform for the LEO Edge” in Proceedings of the 4th International Workshop on Edge Systems, Analytics and Networking (EdgeSys 2021), Apr. 2021, pp. 43–48.




Easy Enough?






We'll need a compute platform



Mobility & Coverage: Virtual Stationarity



Mobility & Coverage: Virtual Stationarity



[2] D. Bhattacharjee, S. Kassing, M. Licciardello, and A. Singla, “In-orbit computing: An outlandish thought experiment?” in Proceedings of the 19th ACM Workshop on Hot Topics in Networks (HotNets 2019), Nov. 2020, pp. 197–204.

[3] T. Pfandzelter et al., “Edge (of the Earth) Replication: Optimizing Content Delivery in Large LEO Satellite Communication Networks” in Proceedings of the 21st IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGrid 2021), May. 2021, pp. 565–575.


Scalability & Multi-Tenancy

Starlink Phase I:

- 4,409 satellites in five “shells”
- Walker Delta constellation
- Each shell is independent


Altitude in km	Planes	Nodes per Plane	Orbital Inclination	Color
550	72	22	53°	blue
1,110	32	50	53.8°	orange
1,130	8	50	74°	purple
1,275	5	75	81°	red
1,325	6	75	70°	green

[18] G. Karabulut Kurt, “An Overview of Communications and Networking in LEO Mega-Constellations,” Working in Space Together: IEEE Low-Earth Orbit Satellite Systems Workshop, Nov. 2022.



Scalability & Multi-Tenancy: Krios

- Scaling services across lots of servers? Kubernetes is always the answer!!!
- K8s is “planet scale”, but we’re going beyond 😊



Satellite handover leads to downtime

Figure 3: Downtime for the four scenarios - default Kubernetes, pessimistic Kubernetes, Kubernetes with just path awareness and Krios handoff.

[1] V. Bhosale, K. Bhardwaj, and A. Gavrilovska, “Toward loosely coupled orchestration for the LEO satellite edge,” in Proceedings of the 3rd USENIX Workshop on Hot Topics in Edge Computing (HotEdge ‘20), Jun. 2020.

Scalability & Multi-Tenancy: Krios

- *Krios*: K8s scheduler that incorporates satellite mobility model, proactively migrates services




Figure 3: Downtime for the four scenarios - default Kubernetes, pessimistic Kubernetes, Kubernetes with just path awareness and Krios handoff.




Figure 2: High level design of Krios components - geographically(left) and the network connections (right).

Proactive support for satellite mobility leads to always-on service


- [1] V. Bhosale, K. Bhardwaj, and A. Gavrilovska, "Toward loosely coupled orchestration for the LEO satellite edge," in Proceedings of the 3rd USENIX Workshop on Hot Topics in Edge Computing (HotEdge '20), Jun. 2020.

Issues with K8s-style orchestration

- Low soft-error rate per server (10^{-3} to 10^{-4} per day) → 1-10 faults every 2 days for 5,000 satellites
→ expensive recovery for containers
- Container migration is very expensive (checkpoint + restore)
→ either downtime or synchronization (development effort)
- Resource constraints
→ containers block resources even when service is not used
→ Alternative: over-provisioning (better hope not too many people use the services) and isolation is not an issue...
- Scalability + elasticity are expensive and high-latency

[5] T. Pfandzelter, J. Hasenburg, and D. Bermbach, “Towards a Computing Platform for the LEO Edge” in Proceedings of the 4th International Workshop on Edge Systems, Analytics and Networking (EdgeSys ‘21), Apr. 2021, pp. 43–48.

Serverless + FaaS



More control given to platform

[5] T. Pfandzelter, J. Hasenburg, and D. Bermbach, "Towards a Computing Platform for the LEO Edge" in Proceedings of the 4th International Workshop on Edge Systems, Analytics and Networking (EdgeSys '21), Apr. 2021, pp. 43–48.


Serverless + FaaS

- De-coupling state (data management platform) and services (FaaS platform)
 - Migration is cheap (data can be replicated online)
 - Scalability is cheap (across satellite servers and within the same server)
 - Fault-tolerance is cheaper (if state copies are kept)
 - Lower resource footprint


[5] T. Pfandzelter, J. Hasenburg, and D. Bermbach, “Towards a Computing Platform for the LEO Edge” in Proceedings of the 4th International Workshop on Edge Systems, Analytics and Networking (EdgeSys ‘21), Apr. 2021, pp. 43–48.

Testing Software Abstractions


We'll need a compute platform



We'll need a compute platform



We'll need a compute platform



CELESTIAL


- Run **cheap** testbeds of LEO edge constellations in the cloud
- Provide isolated computing resources for each satellite in a scalable manner
- Support arbitrary software systems, including platforms
- Reliable efficiency, accuracy, reproducibility, repeatability

[18] T. Pfandzelter, and D. Bermbach, “Celestial: Virtual Software System Testbeds for the LEO Edge” in *Proceedings of the 23rd ACM/IFIP International Middleware Conference (Middleware ’22)*, Nov. 2022.

CELESTIAL

- Run **cheap** testbeds of LEO edge constellations in the cloud
- Provide **isolated computing resources for each satellite** in a scalable manner
- Support arbitrary software systems, including platforms
- Reliable efficiency, accuracy, reproducibility, repeatability

This seems hard!



[18] T. Pfandzelter, and D. Bermbach, “Celestial: Virtual Software System Testbeds for the LEO Edge” in *Proceedings of the 23rd ACM/IFIP International Middleware Conference (Middleware ‘22)*, Nov. 2022.

Thousands of Isolated Compute Resources in the Cloud vs. “Cheap”

	VMs		Containers
Cost	High		Low
Isolation	High		Low
Software Support	High		Medium
Overhead	High		Low

[18] T. Pfandzelter, and D. Bermbach, “Celestial: Virtual Software System Testbeds for the LEO Edge” in *Proceedings of the 23rd ACM/IFIP International Middleware Conference (Middleware ’22)*, Nov. 2022.

Thousands of Isolated Compute Resources in the Cloud vs. “Cheap”

	VMs	MicroVMs	Containers
Cost	High	Low	Low
Isolation	High	High	Low
Software Support	High	High	Medium
Overhead	High	Low	Low


[5] A. Agache, M. Brooker, A. Iordache, A. Liguori, R. Neugebauer, P. Piwonka, and D.-M. Popa, “Firecracker: Lightweight Virtualization for Serverless Applications” in Proceedings of the 17th USENIX Symposium on Networked Systems Design and Implementation (NSDI ’20), Feb. 2022, pp. 419–434.

1. Efficient virtualization with microVMs

2. Bounding box


3. Scalable network emulation

Thousands?




[2] D. Bhattacharjee, S. Kassing, M. Licciardello, and A. Singla, “In-orbit computing: An outlandish thought experiment?” in Proceedings of the 19th ACM Workshop on Hot Topics in Networks (HotNets 2019), Nov. 2020, pp. 197–204.


Thousands?



Bounding Box




Bounding Box



1. Efficient virtualization with microVMs


2. Bounding box

3. Scalable network emulation




[13] S. Becker et al., “Network Emulation in Large-Scale Virtual Edge Testbeds: A Note of Caution and the Way Forward ” in Proceedings of the 2nd International Workshop on Testing Distributed Internet of Things Systems (TDIS 2022), Sep. 2022.

Scalable Network Emulation




[13] S. Becker et al., “Network Emulation in Large-Scale Virtual Edge Testbeds: A Note of Caution and the Way Forward ” in Proceedings of the 2nd International Workshop on Testing Distributed Internet of Things Systems (TDIS 2022), Sep. 2022.




[13] S. Becker et al., “Network Emulation in Large-Scale Virtual Edge Testbeds: A Note of Caution and the Way Forward ” in Proceedings of the 2nd International Workshop on Testing Distributed Internet of Things Systems (TDIS 2022), Sep. 2022.

1. Efficient virtualization with microVMs
2. Bounding box
3. Scalable network emulation


Celestial: Overview




Celestial: Overview




Celestial: Overview




Celestial: Overview



Example Deployment



LEO Tooling Ecosystem



Deploying

Where do we allocate servers?*

*same goes for static services (database replicas), etc.

Where do we allocate servers?

At every satellite?

Where do we allocate servers?

~~At every satellite?~~ Too expensive!

Where do we allocate servers?

~~At every satellite?~~ Too expensive!

Where demand is?

Where do we allocate servers?

- ~~At every satellite?~~ Too expensive!
- ~~Where demand is?~~ Impossible (mobility)

Where do we allocate servers?

- ~~At every satellite?~~ Too expensive!
- ~~Where demand is?~~ Impossible (mobility)

Evenly across the constellation?


Where do we allocate servers?

~~At every satellite?~~ Too expensive!


~~Where demand is?~~ Impossible (mobility)

Evenly across the constellation? How?


LEO Satellite Network Topology




LEO Satellite Network Topology





LEO Satellite Network Topology





LEO Satellite Network Topology





LEO Satellite Network Topology





LEO Satellite Network Topology




LEO Satellite Network Topology




LEO Satellite Network Topology



LEO Satellite Network Topology



= 2D Torus



Resource Placement on 2D Tori

- Well-studied problem
- Had some grid/cluster computing applications in the 1990s
- Considers the discrete „hop“ distance metric
- Bae [27], Alb Daiwi and Livingston [28], Bae and Bose [29] propose solutions based on Lee error correcting codes
- „Perfect“ solutions exist under some constraints
- Non-perfect solutions exist for all tori




Figure 1. A perfect distance-2 placement in a 13×13 torus (wraparound edges not shown).

[21]

[21] M. M. Bae, “Resource placement in torus-based networks,” IEEE Transactions on Computers, vol. 46, no. 10, pp. 1083–1092, 1997.

[22] B. F. Alb Daiwi and M. L. Livingston, “Perfect distance-d placements in 2D toroidal networks,” Journal of Supercomputing, vol. 29, no. 1, pp. 45–57, 2004.

[23] M. M. Bae and B. Bose, “Spare processor allocation for fault tolerance in torus-based multicomputers,” in Proceedings of the Annual Symposium on Fault Tolerant Computing, Jun. 1996, pp. 282–291.

Resource Placement on 2D Tori

- Well-studied problem
- Had some grid/cluster computing applications in the 1990s
- Considers the **discrete „hop“ distance metric**
- Bae [27], Albdaoui and Livingston [28], Bae and Bose [29] propose solutions based on Lee error correcting codes
- „Perfect“ solutions exist under some constraints
- Non-perfect solutions exist for all tori




Figure 1. A perfect distance-2 placement in a 13×13 torus (wraparound edges not shown).


[21]

[21] M. M. Bae, “Resource placement in torus-based networks,” IEEE Transactions on Computers, vol. 46, no. 10, pp. 1083–1092, 1997.


[22] B. F. Albdaoui and M. L. Livingston, “Perfect distance-d placements in 2D toroidal networks,” Journal of Supercomputing, vol. 29, no. 1, pp. 45–57, 2004.

[23] M. M. Bae and B. Bose, “Spare processor allocation for fault tolerance in torus-based multicomputers,” in Proceedings of the Annual Symposium on Fault Tolerant Computing, Jun. 1996, pp. 282–291.

The Issue with Discrete Distances: Two Starlink Shells




Planes (N): **8**
Nodes/Plane (M): **50**
Altitude: **1,130km**
Inclination: **74°**




Planes (M): **72**
Nodes/Plane (M): **22**
Altitude: **550km**
Inclination: **53°**

The Issue with Discrete Distances: Two Starlink Shells



Planes (N): **8**
Nodes/Plane (M): **50**
Altitude: **1,130km**
Inclination: **74°**




Planes (M): **72**
Nodes/Plane (M): **22**
Altitude: **550km**
Inclination: **53°**

Calculating Distances

- D_M is easy as satellites are evenly spaced along their orbit

$$D_M = (\underbrace{r_E}_{\text{Earth Radius}} + \underbrace{h}_{\text{Altitude}}) \sqrt{2(1 - \cos(\frac{2\pi}{M}))}$$

* assuming the Earth is a sphere



[24] T. Pfandzelter et al., "QoS-Aware Resource Placement for LEO Satellite Edge Computing," in Proceedings of the 6th IEEE International Conference on Fog and Edge Computing (ICFEC 2022), May 2022, pp. 66–72.

Calculating Distances

- D_N changes over time
- Can only use aggregate distances
- *Maximum* distance gives upper limit


$$D_N^{\max} = \max_{t \in [0, T]} D_N = (r_E + h) \sqrt{2(1 - \cos(\frac{2\pi}{N}))}$$

- *Mean* distance gives a closer approximation of effects observed over time

$$\begin{aligned} \bar{D}_N &= \frac{1}{T} \int_0^T D_N(t) dx \\ &= \frac{2}{\pi} (r_E + h) \sqrt{2(1 - \cos(\frac{2\pi}{N}))} E(1 - \cos^2(i)) \end{aligned}$$

* assuming the Earth is a sphere


[24] T. Pfandzelter et al., "QoS-Aware Resource Placement for LEO Satellite Edge Computing," in Proceedings of the 6th IEEE International Conference on Fog and Edge Computing (ICFEC 2022), May 2022, pp. 66–72.



Integrating Real Distances

Example: 5x3 torus

- $N = 5$
- $M = 3$
- $D_N = 0.8$
- $D_M = 1.3$
- Target distance $d = 1.4$



[24] T. Pfandzelter et al., "QoS-Aware Resource Placement for LEO Satellite Edge Computing," in Proceedings of the 6th IEEE International Conference on Fog and Edge Computing (ICFEC 2022), May 2022, pp. 66–72.


Integrating Real Distances

- Introduce **virtual nodes** so that:

$$M' = \left\lceil M \frac{D_M}{D_N} \right\rceil$$

- Here:

$$M' = \left\lceil 3 \frac{1.3}{0.8} \right\rceil = 5$$




[24] T. Pfandzelter et al., "QoS-Aware Resource Placement for LEO Satellite Edge Computing," in Proceedings of the 6th IEEE International Conference on Fog and Edge Computing (ICFEC 2022), May 2022, pp. 66–72.

Integrating Real Distances

- Perform a $\left\lfloor \frac{d}{D_M} \right\rfloor$ -hops placement using the known algorithms
- Here:

$$\left\lfloor \frac{d}{D_M} \right\rfloor = \left\lfloor \frac{1.4}{1.3} \right\rfloor = 1$$


, so we do a 1-hop placement



[24] T. Pfandzelter et al., “QoS-Aware Resource Placement for LEO Satellite Edge Computing,” in Proceedings of the 6th IEEE International Conference on Fog and Edge Computing (ICFEC 2022), May 2022, pp. 66–72.

Integrating Real Distances


- Multiply y-coordinates of resource nodes with $\frac{M}{M'}$ to have the placement on the original torus




[24] T. Pfandzelter et al., “QoS-Aware Resource Placement for LEO Satellite Edge Computing,” in Proceedings of the 6th IEEE International Conference on Fog and Edge Computing (ICFEC 2022), May 2022, pp. 66–72.

Other Orbital Edge Computing

Types of Edge Computing in Orbit




- In-network computing
- Piggy-backing LEO satellite broadband networks
- Same applications as terrestrial fog computing
- Large scale, highly dynamic




- Edge, on-device computing
- Processing to lower downlink bandwidth and energy consumption
- On-satellite sensor data (e.g., imaging)
- (Mostly) isolated

Computers in Space: HPE Spaceborne Computer



[11] Hewlett Packard Enterprise, “HPE Spaceborne Computer”, Online:
<https://www.hpe.com/us/en/compute/hpc/supercomputing/spaceborne.html>

Computers in Space: Tiansuan Constellation



[7] S. Wang, “Tiansuan Constellation: An Open Research Platform for Satellite Network and Computing” in LEOCONN 2022, Oct. 2022.


Computers in Space: D-Orbit



[25] D-Orbit, “D-Orbit Cloud Platform”, Online: <https://www.dorbit.space/advanced-services>

Edge Computing on Nano Satellites


- Nanosatellites are much cheaper to produce and launch (e.g., 10,000x cheaper), very common now
- Size limits power, limits compute and storage capabilities
→ How do you decide which data to store between access to ground station?
- (Collaboratively) processing data on the satellite helps filter uninteresting data




[9] B. Denby, and B. Lucia, “Orbital edge computing: Nanosatellite constellations as a new class of computer system,” in Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2020), Mar. 2020, pp. 939–954.

Mars

Martian Edge Computing



Martian Edge Computing – Analogy: Mars Sample Return




Landing stuff on Mars is expensive and difficult!

[26] A. Witze, "NASA plans Mars sample-return rover", Nature 509, pp. 272, May 2014.

Martian Edge Computing

- Sending servers to space is expensive → why do we need to land them?
- Land cheap sensors, don't care if they break
- Keep servers in Medium- to Low-Mars Orbit and share them between ground-sensors
- Relay satellite for communication to Earth



Summary

Summary

- Computing in Space uses Fog Computing techniques
- Difference between:
 1. Fog Computing transferred to LEO Networks
 2. In-Orbit Edge Computing
- LEO Network Fog Computing has additional challenges to terrestrial fog computing
 - Now your servers are moving!
 - Less resources!
 - Larger scale
 - More fault tolerance necessary (can't send an engineer to fix your servers)

Yes, there are projects/theses 😊

Questions?



References (I)

- [1] V. Bhosale, K. Bhardwaj, and A. Gavrilovska, “Toward loosely coupled orchestration for the LEO satellite edge,” in *Proceedings of the 3rd USENIX Workshop on Hot Topics in Edge Computing (HotEdge '20)*, Jun. 2020.
- [2] D. Bhattacharjee, S. Kassing, M. Licciardello, and A. Singla, “In-orbit computing: An outlandish thought experiment?” in *Proceedings of the 19th ACM Workshop on Hot Topics in Networks (HotNets '19)*, Nov. 2020, pp. 197–204.
- [3] T. Pfandzelter, and D. Bermbach, “Edge (of the Earth) Replication: Optimizing Content Delivery in Large LEO Satellite Communication Networks” in *Proceedings of the 21st IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGrid 2021)*, May. 2021, pp. 565–575.
- [4] U. Speidel, “LEO mega-constellations - where the challenges lie” in *LEOCONN 2022*, Oct. 2022.
- [5] T. Pfandzelter, J. Hasenburg, and D. Bermbach, “Towards a Computing Platform for the LEO Edge” in *Proceedings of the 4th International Workshop on Edge Systems, Analytics and Networking (EdgeSys '21)*, Apr. 2021, pp. 43–48.
- [6] A. Agache, M. Brooker, A. Iordache, A. Liguori, R. Neugebauer, P. Piwonka, and D.-M. Popa, “Firecracker: Lightweight Virtualization for Serverless Applications” in *Proceedings of the 17th USENIX Symposium on Networked Systems Design and Implementation (NSDI '20)*, Feb. 2022, pp. 419–434.
- [7] S. Wang, “Tiansuan Constellation: An Open Research Platform for Satellite Network and Computing” in *LEOCONN 2022*, Oct. 2022.
- [8] M. Kassem, “Large-Scale Realistic LEO Networks Emulation” in *LEOCONN 2022*, Oct. 2022.
- [9] B. Denby, and B. Lucia, “Orbital edge computing: Nanosatellite constellations as a new class of computer system,” in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2020)*, Mar. 2020, pp. 939–954.

References (II)


- [10] S. Kassing, D. Bhattacharjee, A. Águas, J. Saethre, and A. Singla, “Exploring the “Internet from space” with Hypatia,” in Proceedings of the ACM Internet Measurement Conference, 2020, pp. 214–229.
- [11] Hewlett Packard Enterprise, “HPE Spaceborne Computer”, Online: <https://www.hpe.com/us/en/compute/hpc/supercomputing/spaceborne.html>
- [12] NOAA Center for Tsunami Research, “DART 4G: Deep-ocean Assessment and Reporting of Tsunami - 4th Generation Tsunami Measurement System,” May. 2016.
- [13] S. Becker, T. Pfandzelter, N. Japke, D. Bermbach, and O. Kao, “Network Emulation in Large-Scale Virtual Edge Testbeds: A Note of Caution and the Way Forward ” in Proceedings of the 2nd International Workshop on Testing Distributed Internet of Things Systems (TDIS 2022), Sep. 2022.
- [14] P. Kirstein, “University College London ARPANET Project, Annual Report 1977,” Apr. 1978.
- [15] S. Wang, Q. Li, M. Xu, X. Ma, A. Zhou, and Q. Sun, “Tiansuan Constellation: An Open Research Platform,” in Proceedings of the 2021 IEEE International Conference on Edge Computing (EDGE), Sep. 2021
- [16] D. Bhattacharjee and A. Singla, “Network topology design at 27,000 km/hour” in Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies (CoNEXT ‘19), Sep. 2022, pp. 341-354.
- [17] T. Pfandzelter and D. Bermbach, “Failure is not an Option: Considerations for Software Fault-Tolerance in LEO Satellite Edge Computing”, Feb. 2023, arXiv:2302.08952.
- [18] T. Pfandzelter, and D. Bermbach, “Celestial: Virtual Software System Testbeds for the LEO Edge” in Proceedings of the 23rd ACM/IFIP International Middleware Conference (Middleware ‘22), Nov. 2022.

References (III)


- [17] T. Pfandzelter and D. Bermbach, “Failure is not an Option: Considerations for Software Fault-Tolerance in LEO Satellite Edge Computing”, Feb. 2023, arXiv:2302.08952.
- [18] G. Karabulut Kurt, “An Overview of Communications and Networking in LEO Mega-Constellations,” Working in Space Together: IEEE Low-Earth Orbit Satellite Systems Workshop, Nov. 2022.
- [19] Akamai Technologies, “Object delivery,” Online: https://developer.akamai.com/legacy/learn/Content_Delivery/Object_Delivery.html 2019.
- [20] A. Akhavain, “Non Terrestrial Networks: Laying out the foundation for ubiquitous network coverage” in LEOCONN Webinar Series, Apr. 2023.
- [21] M. M. Bae, “Resource placement in torus-based networks,” IEEE Transactions on Computers, vol. 46, no. 10, pp. 1083–1092, 1997.
- [22] B. F. Albdaiwi and M. L. Livingston, “Perfect distance-d placements in 2D toroidal networks,” Journal of Supercomputing, vol. 29, no. 1, pp. 45–57, 2004.
- [23] M. M. Bae and B. Bose, “Spare processor allocation for fault tolerance in torus-based multicomputers,” in Proceedings of the Annual Symposium on Fault Tolerant Computing, Jun. 1996, pp. 282–291.
- [24] T. Pfandzelter and D. Bermbach, “QoS-Aware Resource Placement for LEO Satellite Edge Computing,” in Proceedings of the 6th IEEE International Conference on Fog and Edge Computing (ICFEC 2022), May 2022, pp. 66–72.
- [25] D-Orbit, “D-Orbit Cloud Platform”, Online: <https://www.dorbit.space/advanced-services>
- [26] A. Witze, “NASA plans Mars sample-return rover”, Nature 509, pp. 272, May 2014.

Backup

Case Study 1: Multi-User Interaction




Case Study 1: Multi-User Interaction




Stream Latency between Accra and Abuja

Case Study 1: Multi-User Interaction




Measured and Expected Network Latency Between Two Users

Case Study 1: Multi-User Interaction




Reproducibility

Case Study 1: Multi-User Interaction



Host CPU Usage

Case Study 1: Multi-User Interaction




Host Memory Usage

Case Study 2: DART




A Cascadia subduction rupture scenario with a 4G network (black triangles) envisioned close to the trench along the length of potential rupture. Time of tsunami detection by nearest 4G is expected to be on order of 10 minutes.

Three currently deployed DART II systems are identified by bright green triangles.




[12] NOAA Center for Tsunami Research, “DART 4G: Deep-ocean Assessment and Reporting of Tsunami - 4th Generation Tsunami Measurement System,” May. 2016.

Case Study 2: DART




[12] NOAA Center for Tsunami Research, "DART 4G: Deep-ocean Assessment and Reporting of Tsunami - 4th Generation Tsunami Measurement System," May. 2016.


Case Study 2: DART



Case Study 2: DART



Case Study 2: DART



Central Server

Satellite Servers

Sensor Reading Notification Latency